What the Swarm Trades Away

On the geometry of collective stability — March 2026

Waldemar Lis — The Self-Experiencing Universe

archive.org/details/@not_foryou

Closing an Open Question

The two previous simulations in this series placed three systems — the individual honeybee, the honeybee swarm, and the cuttlefish Sepia officinalis — at specific positions in a three-dimensional space describing forms of consciousness. Two of those positions were derived from published experimental data. The third — the swarm — was placed by estimate.

This simulation closes that gap. Using Thomas Seeley's published research on swarm decision-making, thermal regulation, and structural coherence, it derives the swarm's position from data rather than intuition. The result is both more precise and more surprising than the estimate suggested.

The Swarm Is More Stable Than We Thought

The stability axis measures how coherently a system maintains its pattern of relationships over time. For the swarm, three independent mechanisms contribute to this: its ability to hold a constant internal temperature across a wide range of external conditions, its structural integrity under physical stress, and the near-irreversibility of its collective decisions once a threshold is crossed.

Each of these mechanisms scores very high. The swarm maintains its core temperature within a narrow range regardless of whether the outside world is cold or warm. It repairs its hanging cluster shape within seconds when disturbed. And once a decision is made — once the swarm has chosen a new home — it almost never reverses course. Out of thirty observed swarm decisions in Seeley's long-term studies, fewer than one resulted in reversal after the commitment threshold was reached.

Taken together, these three indicators place the swarm's stability at 0.94 — higher than the working estimate of 0.85, and higher than any other system in the current model.

The Swarm Is Far Less Integrated Than We Thought

The integration axis measures how tightly the parts of a system are coupled to each other — whether the system functions as a whole rather than a collection of loosely connected components. This is where the empirical data delivers its most striking result.

The fundamental constraint is time. When neurons communicate within a brain, signals travel in milliseconds. When bees communicate within a swarm — through the waggle dance, through chemical signals, through physical contact — the relevant cycles take seconds to minutes. This gap, spanning roughly five orders of magnitude, structurally limits how tightly integrated the swarm can be as a system.

The quorum mechanism confirms this. A swarm of several hundred foragers reaches a collective decision when approximately fifteen scouts converge on the same site. Fifteen out of four hundred is a small fraction. The system is sensitive to this fraction — the decision cannot happen without it — but the coupling across the full population is sparse rather than dense.

The derived integration value for the swarm is 0.18 — substantially lower than the working estimate of 0.45. The swarm is less than half as integrated as previously assumed.

A Trade-Off That Is Now Measurable

Taken together, the empirical derivation reveals a sharp trade-off. The swarm achieves extraordinary stability — nearly perfect persistence of pattern over time — but at the cost of low integration. Its parts are loosely coupled. Its coherence is not the coherence of a brain, where every region is in constant rapid dialogue with every other. It is the coherence of a crystal: stable because rigid, not because interconnected.

This trade-off was assumed in the previous simulations. It is now measured.

The consequence for the geometry of the model is that the swarm has moved. The distance between the individual bee and the swarm has grown from 0.47 to 0.64. The distance between the swarm and the cuttlefish has grown from 0.79 to 0.95. The swarm is now the most geometrically isolated system in the model — furthest from both the individual bee and from Sepia — not because it is less sophisticated, but because its form of organization occupies a structurally distinct region of the space.

The Central Finding Holds — and Strengthens

The original finding — that an individual honeybee is geometrically closer to a cuttlefish than to its own swarm — was derived from estimates. It now holds with empirical values, and the gap has doubled. The individual bee sits at distance 0.32 from Sepia and distance 0.64 from the swarm it belongs to.

The swarm is not a scaled-up bee. It is not a step up the ladder of consciousness. It is a different shape entirely — one that sacrifices the self-modeling capacity of its individual members in exchange for a stability that no individual could achieve alone.

Status

This is the third simulation in the VCP empirical series. The three-vector model is now fully grounded for the swarm. The individual bee vector and the Sepia vector retain partial empirical support — their s and i values remain estimates pending future empirical grounding. Full technical documentation is available to researchers through archive.org/details/@not_foryou.